

B. Justification and requirement for continued SUMEX use

The research extending VM depends critically upon continued SUMEX support. This support is of two kinds: computer resources and community interaction. Exploiting the VM knowledge requires use of the SUMEX computer resource for several purposes. First, there is the need to continue run the VM program itself to test new rules, to identify the interaction of old and new rules, and to validate new kinds of rules. VM was built to run on SUMEX in INTERLISP, and its use depends upon use of SUMEX. In addition, the knowledge refinement process requires occasional use of SUMEX facilities such as MLAB, to test new mathematical models, and SPSS to analyze the results of studies. Finally, the SUMEX text processing support is tremendously useful.

Continuing as a part of the SUMEX community is particularly important intellectually. The general theme of AI -- representation and manipulation of knowledge -- lies at the heart of the research we conduct on interpretation of medical data. The exploitation of that knowledge involves both basic research into the nature of that knowledge and more applied research to refine and exploit it. It is extremely important that we continue to participate in the research of the AI community in representation and manipulation of knowledge.

D. Recommendations for Future Community and Resource Development

We perceive the evolution of our AI capability as moving from a highly speculative development state, for which the interactive development capabilities of SUMEX are vital, to a more stable but still changing validation-and-evaluation state. Ultimately we foresee rather stable specification of a program for routine clinical use. Thus, we see the need to transfer our AI techniques from the SUMEX PDP-10 to a local host. For this transfer, a principal long-range need is for software systems that will allow us to run AI systems on a mini-computer after they have been developed on the more powerful SUMEX facility. If the validation of PUFF-VM in the PMC clinical setting shows the programs to be effective in health care, then we hope and expect to be able to provide the capability on a routine basis.

We would also like to encourage SUMEX's role as a facilitator of information transfer between AIM users. This can happen by scheduling on-line demonstrations that any other user can "connect to," or by providing a common depository for AI and medicine information. This might take the form of on-line bibliographies, collecting common user packages, or connecting common research interests together. This communication service would complement the technical service facilities currently provided by the SUMEX staff.

II.A.2.5 Rutgers Computers in Biomedicine Project [Rutgers-AIM]

Rutgers Computers in Biomedicine

Rutgers Research Resource--Computers in Biomedicine

Principal Investigators: Saul Amarel and C. A. Kulikowski
Rutgers University
New Brunswick, New Jersey

I. SUMMARY OF RESEARCH PROGRAM

A. Goals and Approach

The fundamental objective of the Rutgers Resource is to develop a computer based framework for significant research in the biomedical sciences and for the application of research results to the solution of important problems in health care. The focal concept is to introduce advanced methods of computer science - particularly in artificial intelligence into specific areas of biomedical inquiry. The computer is used as an integral part of the inquiry process, both for the development and organization of knowledge in a domain and for its utilization in problem solving and in processes of experimentation and theory formation.

At present, the total number of investigators who participate in scientific activities of the Resource is 97, of these, 24 have Rutgers appointments, 25 are outside investigators who participate in collaborative research projects that are mainly located at Rutgers, and 48 are investigators from collaborative national AIM projects that are located in different parts of the country. In addition, the Resource has 11 other members in Administrative, Computer Systems/Operations and general programming and secretarial functions. Thus, the Rutgers Resource community numbers at present a total of 108 participants.

Resource activities include research projects (collaborative research and core research) training/dissemination projects, and computing services in support of user projects.

B. Medical Relevance and Collaborations

In 1981-82 we continued the development of several versatile systems for building and testing consultation models in biomedicine. The EXPERT system has had many of its capabilities enhanced in the course of collaborative research in the areas of rheumatology, ophthalmology, and clinical pathology.

In rheumatology, our collaboration with Drs. Donald Lindberg and Gordon Sharp at the University of Missouri-Columbia has continued at a very active level. The model for rheumatological diseases was expanded to include detailed diagnostic criteria for 26 major diseases, and management

advice/treatment planning was begun for several of them. Dr. Sharp's group continues to develop the knowledge base in this area, with formalization of the knowledge carried out in conjunction with Dr. Lindberg's group and the Medical Expert Systems Group at Rutgers. The Resource researchers have developed new representational elements for EXPERT in response to the needs of the rheumatology research, and Politakis has developed a coordinated system called SEEK (System for Empirical Experimentation with Expert Knowledge) which provides interactive assistance to the human expert in testing, refining and updating a knowledge base against a data base of trial cases.

In treatment planning, the system for treatment selection and planning developed by Kastner has been used in ophthalmology (infectious eye diseases), where Dr. Chandler Dawson of the Proctor Foundation, University of California at San Francisco has developed an expert reasoning model.

In clinical pathology our main collaboration has been with Dr. Robert Galen (Overlook Hospital and Columbia University), where we continued development of the serum protein electrophoresis model so that it was usable clinically and incorporated into an instrument - the scanning densitometer manufactured by Helena Laboratories. The new instrument with interpretive reporting capabilities is now on the market, and represents the first known spin-off of AI expert systems research in the field of laboratory instrumentation.

In biomedical modeling applications we continued to experiment with prototype models in conjunction with Dr. David Garfinkel (enzyme kinetics) and Dr. William Yamamoto (respiration models).

C. Highlights of Research Progress

1. Expert Medical Systems (C. Kulikowski, S. Weiss)

Research has continued on problems of representation, inference and control in expert systems. More emphasis has been placed this year on problems of knowledge base acquisition, empirical testing and refinement of reasoning (the SEEK system), and treatment planning strategies over time. From a technological point of view the market availability of the interpretive reporting version of a scanning densitometer represents an important achievement for AIM research in showing a practical impact in medicine.

1.1 SEEK: A System for Empirical Experimentation with Expert Knowledge

SEEK is a system which has been developed to give interactive advice about rule refinement during the design of an expert system. The advice takes the form of suggestions for possible experiments in generalizing and specializing rules in an expert model that has been specified based on reasoning rules cited by a human expert. Case experience, in the form of stored cases with known conclusions, is used to interactively guide the

expert in refining the rules of a model. The design framework of SEEK consists of a tabular model for expressing expert-modeled rules and a general consultation system for applying a model to specific cases. This approach has proven particularly valuable in assisting the expert in domains where the logic for discriminating two diagnoses is difficult to specify; and we have benefitted primarily from experience in building the consultation system in rheumatology.

1.2 Treatment Planning

The ranking and selection strategies developed as a stand-alone system last year have been incorporated into the EXPERT framework. New capabilities for expressing reasoning over time have been added, so stored chart reviews can be carried out automatically, summarizing various patterns of findings over time, and abstracting the major features of interest for prognostic advice or treatment recommendations. Applications have been in infectious eye disease modeling, rheumatology treatment, and sequential advice in interpretation and sequencing of cardiac enzyme tests (e.g. CPK/LDH isoenzymes).

1.3 Technology Transfer

The automated translation of an EXPERT model with some representational constraints into an algorithmic model which is then also translatable automatically into a microprocessor assembly language has been demonstrated by producing the serum protein electrophoresis model for interpretation.

1.4 Representational Extensions to EXPERT

We have created a number of special front-end and reporting facilities to make it easier to enter data, and to flexibly format advice at the output.

We have developed an entirely new representation for expressing time dependencies of actions that are first recommended and then monitored for matching over time.

We have developed a new representation for creating a special class of instantiations in the conclusions of EXPERT without sacrificing the efficiency of model precompilation.

1.5 Learning with Prior Structural Knowledge

We have developed a multiple knowledge source/blackboard model for learning with prior knowledge, and applied it in the area of glaucoma and thyroid modeling.

2. Models of Planning and Commonsense Reasoning (C. Schmidt, N.S. Sridharan)

2.1 Plan Generation

In the last year, we have incorporated several features into the plan generation process that enables it to deal with incomplete knowledge of situations. Our planning programs include explicit annotations of decisions taken and assumptions made. These annotations, coupled with a structured representation of the planning process, provide a basis for plan revision when new information becomes available. Our plan generator has now the capacity to construct and manipulate descriptions of objects required in the plan even though knowledge of the situation may not include explicit information about the existence of these objects.

2.2 Reasoning with Commonsense Concepts

In this work we assume that (i) concepts are structured hierarchically and (ii) associated with a concept are intentional descriptions which define either necessary or typical aspects of individuals that are members of the concept set. Proceeding within the framework of default reasoning developed by Reiter, we have defined and implemented a system which uses the concept hierarchy and associated intentional descriptions to systematically generate default theories from which default inferences about an individual may be drawn.

3. Artificial Intelligence: Representations, Reasoning, System Development

3.1 Representations and Inference in AIMDS (N.S. Sridharan)

The two main applications of AIMDS continue to be Psychological models for Commonsense Reasoning and Cognitive models of Legal Argumentation. In addition to work in these application areas, we introduced recently new features to the system, and we made progress in performance enhancements and in improved interfaces to the user. The new features include refinements of the formalism of Descriptions, mechanisms for representing and using structured definitions, and the introduction of a limited facility for arithmetic functions. We devoted considerable effort to the development of short-term and long-term solutions to questions of improved system performance. AIMDS was adapted to run with Elisp, to give us more address space. Improvements in user interface included a facility for loading a knowledge base in pieces from a collection of external files, and an interactive facility for composition of Descriptions.

3.2 Expertise Acquisition (S. Amarel, T. Mitchell)

There are two main research activities in this area. One is focusing on the task of learning problem solving heuristics by experimentation, and the other is concerned with improvements in problem solving expertise via shifts in problem representation, i.e., via reformulation.

We have now completed the initial design, implementation, and partial testing of a heuristic search program, called LEX, that learns heuristics to guide its search in a particular task area: solving symbolic integration problems. More recently, we completed a prototype 'Problem Generator' module, which will allow LEX to generate its own practice problems. We view the current system as a laboratory in which we can now conduct experiments and study specific issues related to machine learning of expertise.

In recent work on changes in Expert Behavior via problem reformulations we developed a conceptual framework for handling problem formulations in which the grammatical specification of solutions for a problem class plays a key role. Within this framework, we conducted theoretical studies of representational shifts in specific classes of problems of reasoning about actions. Relationships between characteristics of 'intelligent expert' systems and developmental processes of problem reformulation were explored; and specific conclusions were obtained regarding the types of knowledge and the nature of mechanisms needed for an expert system to improve its performance with experience.

3.3 Programming Environment for AI Research (C. Hedrick)

During the last year we have been concentrating on two issues: (i) solving the address space problem with the PDP-10 architecture; and (ii) exploratory work with personal computers and networking. We have now completed versions of Rutgers/UCI LISP, PASCAL and SETL that use the full 23-bit address space available on the DEC-2060 processor; and we are currently working on a large-address space implementation of Common LISP for DEC-20. We expect to install in the summer of '82 two Dolphin personal computers, and we developed much of the Ethernet software that we will use to connect the Dolphins with each other and with the DEC-20. SUMEX staff has been providing help in planning for the Dolphins and for the local network.

D. Up-to-Date List of Publications

The following is an update of publications in the Rutgers Resource for the period 1981 and 1982 (only publications not listed in previous SUMEX annual reports are presented here).

Amarel, S. "Problems of Representation in Heuristic Problem Solving; Related Issues in the Development of Expert Systems." Feb. 1981. CBM-TR-118, also to appear in "Methods of Heuristics", Groner, Groner and Bishof (eds), Laurence Erlbaum.*

Amarel, S. "Expert Behavior and Problem Representations." Feb. 1982. CBM-TR-126; also to appear in "Artificial and Human Intelligence," Elithorn and Banerji (eds), 1982.

Bresina, J. "An Interactive Planner that Creates a Structured, Annotated Trace of its Operation." Sept. 1981. CBM-TR-123.

Drastal, G. and Kulikowski, C. "Knowledge Based Acquisition of Rules for Medical Diagnosis", Department of Computer Science Technical Report No. CBM-TM-92, Rutgers University, March 1981.*

Keller, R.M. and Nagel, D.J. (1981), "Some Experiments in Abstraction of Relational Characteristics", Department of Computer Science Technical Report No. DCS-TM-15. Rutgers University, March 1981.*

Kulikowski, C.A. "Computers and Logic: Expert Systems for Thyroid Function Testing", Diagnostic Medicine, 99-102, July 1981.* Kulikowski, C.A., Weiss, S., Galen, R.S. (1981). "Computerized Diagnosis in the Lab", Medical Laboratory Observer: 41-57, July 1981.*

Mitchell, T.M. (1981). "Generalization as Search", Artificial Intelligence, to appear 1981.*

Mitchell, T.M. (1981). "Combining Empirical and Analytical Methods for Inferring Heuristics", Proceedings at the NATO Symposium on Artificial and Human Intelligence, Lyon, France, October 1981, Department of Computer Science Technical Report No. LCSR-TR-21.*

Mitchell, T.M., Utgoff, P.E., Nudel, B. and Banerji, R. (1981). "Learning Problem-Solving Heuristics through Practice", Department of Computer Science Technical Report No. LCSR-TR-15, Rutgers University, July 1981; also in Proceedings at IJCAI 7, Vancouver, BC, August 1981.*

Nudel, B. and Utgoff, P. "A Bibliography on Machine Learning." April 1981. CBM-TR-120.

Schwanke, R. "Common Virtual Arrays in PDP-11 Fortran: An Exercise in Software Engineering." Nov. 1981. CBM-TR-124.

Smith, D. (1981). "A Strategic Interaction Paradigm for Language Acquisition", forthcoming Ph.D. dissertation, Department of Computer Science, Rutgers University, expected December, 1981.*

Sridharan, N.S. "Representing Knowledge in Introduction Using TAXMAN Examples." Nov. 1981. CBM-TR-125/LRP-TR-12.

Sridharan, N.S., Goodson, J.L. and Schmidt, C.F. "A Research Strategy for Computational Studies of Event and Action Perception." July 1981. CBM-TR-122.

Weiss, S., Kern, K., Kulikowski, C. and Ushchold, M. "A Guide to the Use of the EXPERT Consultation System." May 1981. CBM-TR-94.

Weiss, S., Kulikowski, C. Expert Consultation Systems: The EXPERT and CASNET Projects, Machine Intelligence, Pergamon INFOTECH, Series 9, Number 3, pp. 339-352 (1981).*

Weiss, S., Kulikowski, C.A. and Galen, R. "Developing Microprocessor Based Expert Models for Instrument Interpretation." June 1981, CBM-TR-121, also in Proceedings IJCAI 7, Vancouver, BC, August 1981.*

*Indicates that the resource was given credit.

E. Funding Support

The Rutgers Resource is funded through an NIH grant entitled "Rutgers Research Resource on Computers in Biomedicine" -- grant number P41RR643. The Co-Principal Investigators are Dr. Saul Amarel, Professor, Chairman Department of Computer Science, and Director of the Laboratory for Computer Science Research, and Dr. Casimir Kulikowski, Professor of Computer Science, Rutgers, the State University of New Jersey.

This grant is in the second year of its fourth 3-year renewal extending from December 1, 1980 through November 30, 1983. The total direct cost awarded for the 3-year period is \$1,404,075 with second year funding of \$460,944 in direct costs from December 1, 1981 through November 30, 1982.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Dissemination

The SUMEX-AIM facility provides a backup node where some of our medical collaborators can access programs developed at Rutgers. The bulk of the medical collaborative work outlined in I.B. above is centered at the Rutgers facility (the Rutgers-AIM node).

Dissemination activities continue to be an important responsibility of the Rutgers Resource within the AIM community. The following activities took place in the last year:

1) Seventh AIM Workshop (1981):

Organized by C. Kulikowski at the Rutgers Research Resource, it was held at the University of British Columbia in conjunction with IJCAI. It consisted of a very successful one-day of summary presentations by members of the AIM community about work in progress, and an evening dinner meeting at which Dr. Allen Newell presented the keynote address.

2) IJCAI-81:

Presentations were made by several Resource investigators, including Drs. Kulikowski, Weiss, and graduate students on the medical projects (SPE densitometer interpretation, treatment planning); and also by Dr. Sridharan and Dr. Mitchell.

3) Society for Computer Medicine Meeting:

Dr. Kulikowski presented work on expert systems in a panel session at the Washington, DC conference of the Society for Computer Medicine.

4) Hawaii International Conference On Systems Sciences:

Dr. Weiss and George Drastal presented papers on the SEEK system, and the a-priori learning model at the medical sessions of this conference.

5) American Association of Clinical Pathologists:

Dr. Kulikowski presented the work on expert systems, and their application in clinical pathology at this meeting, while Drs. Lindberg and Galen also participated and touched on these topics in tutorials and papers.

6) American Rheumatology Association:

Dr. Sharp and his group presented the rheumatology knowledge base and consultation program at this meeting in a poster session.

7) New York Academy of Medicine:

Dr. Kulikowski presented a paper on the expert systems in medicine at the March '82 meeting of the Academy.

8) Conference of Computers in Health Care, 1981:

The Rutgers Resource was a co-sponsor.

9) NATO Symposium on Human and Artificial Intelligence:

Dr. Amarel and Dr. Mitchell presented papers on Expertise acquisition at this Symposium, which was held in Lyon, France in October 1981.

10) Lectures at Jilin University, China:

Dr. Amarel and Dr. Kulikowski were invited to Jilin University for a series of lectures. Work on medical expert systems and AI was presented in these lectures.

B. National AIM Projects at Rutgers

The national AIM projects, approved by the AIM Executive Committee, that are associated with the Rutgers-AIM node are the following:

1) BRIGHT project, under the direction of Dr. W. Gordon Walker of Johns Hopkins University, mostly for research in clinical medicine.

2) CONGEN project on Computers in Chemistry, under Dr. Djerassi and Dr. D. Smith from Stanford University, are predominantly at SUMEX; they are using the Rutgers facility as a backup for demonstrations.

3) INTERNIST/CADUCEUS project, headed by Dr. Myers and Dr. Pople from the University of Pittsburgh, has been using the Rutgers Resource as a backup system for development and experimentation.

4) Medical Knowledge Representation project, headed by Dr. Chandrasekaran from Ohio State University, is doing most of its research on the Rutgers system.

5) PURSUIT project, directed by Dr. Greenes from Harvard University, is doing most of its research on a Goal-Directed Model of Clinical Decision-Making at Rutgers.

6) Biomedical Modeling, by Dr. Garfinkel from the University of Pennsylvania.

7) MEDSIM project: This is a pilot project designed to provide resource-sharing and community building facilities for about 25 researchers in biomathematical modeling and simulation.

C. Critique of SUMEX-AIM Resource Management

Rutgers is currently using the SUMEX facility primarily for communication with other researchers in the AIM community and with SUMEX staff, and also for backup computing in demonstrations, conferences and site visits. Our usage is currently running at less than 50 connect hours per year at SUMEX, with an overall connect/CPU ratio of about 30.

In addition to the computer usage, we have benefitted extensively from SUMEX staff expertise in the area of local networking and personal computers. The AIM Executive Committee has allocated to the Rutgers-AIM node one of the Xerox Dolphins acquired by SUMEX, so that - together with a second Dolphin provided by DARPA to Rutgers - we can develop experience with the new personal machine environment for research in AI systems. We expect the Dolphins to be installed at Rutgers in summer '82. The SUMEX staff has provided the software that we will be using to implement Ethernet service for the personal computers, as well as advice on the logistics of installing them. As in the past, the quality of support and cooperation given by SUMEX staff has been very high.

III. RESEARCH PLANS

A. Project Goals and Plans

We are planning to continue along the main lines of research that we have established in the Resource to date. Our medical collaborations will continue with emphasis on development of expert consultation systems in rheumatology, ophthalmology and clinical pathology. The basic AI issues of plan recognition/generation and default reasoning will continue to receive attention. Our core work will continue with emphasis on further development of the EXPERT framework and also on AI studies in representations and problems of knowledge and expertise acquisition. We

also plan to continue our participation in AIM dissemination and training activities as well as our contribution -- via the RUGERS/LCSR computer -- to the shared computing facilities of the national AIM network.

B. Justification and Requirements for Continued SUMEX Use

Continued access to SUMEX is needed for:

- 1) Backup for demos, etc.
- 2) Programs developed to serve the National AIM Community should be runnable on both facilities.
- 3) There should be joint development activities between the staffs at Rutgers and SUMEX in order to ensure portability, share the load, and provide a wider variety of inputs for developments.

C. Needs and Plans for Other Computing Resources Beyond SUMEX-AIM

Beyond the current SUMEX-AIM facility there is need for access to a more 'personal' type of computing facilities (e.g., Dolphins).

D. Recommendations for Future Community and Resource Development

We strongly support the efforts of SUMEX staff to develop personal computing resources. Our own local planning efforts have been helped considerably by the assistance of SUMEX. As physical facilities are spread out in more and more places, we believe that centers such as SUMEX can play a growing role in keeping track of technology, planning, and technical support of users. It may also prove convenient for SUMEX to serve as a center for mail and other communications activities carried out by personal computers at remote sites.

Our experience during the last year confirms our opinion that an increasingly important role of the SUMEX staff will be to function as a central consulting and software development organization. They are currently performing this role in regard to personal computers and networking, as well as more traditional roles such as INTERLISP support.

II.A.2.6 SECS - Simulation and Evaluation of Chemical Synthesis

SECS - Simulation and Evaluation of Chemical Synthesis

Principal Investigator: W. Todd Wipke
Board of Studies in Chemistry
University of California
Santa Cruz, CA. 95064

Coworkers:

D. Dolata	(Grad student)
I. Kim	(Grad student)
D. Rogers	(Grad Student)
J. Chou	(Postdoctoral)
P. Condran	(Postdoctoral)
T. Moock	(Postdoctoral)
A. Cabral	(Postdoctoral)

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The long range goal of this project is to develop the logical principles of molecular construction and to use these in developing practical computer programs to assist investigators in designing stereospecific syntheses of complex bio-organic molecules. Our specific goals this past year included conversion of SECS to machine independent code, continued exploration of strategies based on starting materials, developing a reasoning program, and expansion of the ALCHEM language for handling precursor evaluation.

The objectives for the XENO project were to establish extensive collaborations with metabolism experimentalists to test XENO predictions and to begin development on methods for assessing the potential biological activity of each metabolite.

B. Medical Relevance and Collaboration

The development of new drugs and the study of how drug structure is related to biological activity depends upon the chemist's ability to synthesize new molecules as well as his ability to modify existing structures, e.g., incorporating isotopic labels or other substituents into biomolecular substrates. The Simulation and Evaluation of Chemical Synthesis (SECS) project aims at assisting the synthetic chemist in designing stereospecific syntheses of biologically important molecules. The advantages of this computer approach over normal manual approaches are many: 1) greater speed in designing a synthesis; 2) freedom from bias of

past experience and past solutions; 3) thorough consideration of all possible syntheses using a more extensive library of chemical reactions than any individual person can remember; 4) greater capability of the computer to deal with the many structures which result; and 5) capability of computer to see molecules in a graph theoretical sense, free from the bias of 2-D projection.

The objective of using XENO (a spinoff of SECS) in metabolism studies is to predict the plausible metabolites of a given xenobiotic in order that they may be analyzed for possible carcinogenicity. Metabolism research may also find this useful in the identification of metabolites in that it suggests what to look for. Finally, one may envision applications of this technology in problem domains where one wishes to alter molecules in order to inhibit certain types of metabolism.

C. Highlights of Research Progress

1. SECS Program Developments

The Simulation and Evaluation of Chemical Synthesis (SECS) program has undergone many additions to improve its capabilities and usefulness to synthetic chemists. We have added a capacity for user-defined transforms, SECS library augmentation, and for multiple step strategies. In addition we have augmented SECS with the SST and QED programs, the former for selection of starting materials for a synthetic target, and the latter for interpreting and maintaining a rule database of chemical synthetic strategies.

a. Perception and Evaluation of Synthetic Precursors: Extensions to the ALCHEM Language

The ALCHEM language has been extended to handle transforms with two sections, one for the reactants and one for the products. This required complete logic revamping so that the precursors can be perceived and evaluated in the same manner as the original target. Symmetry and steric constraints are now expressible with respect to the precursors. Additionally, the ALCHEM compiler has been given semantic checking capabilities for greater error detection.

b. User Defined Transform / Library Augmentation Module

As part of a concerted effort to improve and maintain the SECS library database, we have implemented the capability for acquisition of chemical knowledge at execution time. This capability complements the existing offline acquisition capability involving translation of detailed transforms written in the chemistry-specific language ALCHEM.

The User Defined Transform module is accessible through both graphical and textual interfaces with the user. The module is designed to accept all information necessary to stereospecifically convert the target molecule to the desired precursor. The reaction specifications may, at the users option, be recorded so that they may be generalized, stored, and

later applied to structurally similar molecules in the current (as well as in future) analysis. Currently the Library Augmentation Module possesses the capability to acquire scope, limitations, and reference information, as well as perform consistency checking to minimize the possibilities for including erroneous information in the database.

It should be noted that the above capability to accept expert knowledge from the chemist at runtime is unique among automated synthesis design programs. Further additions to allow the input of more defined scope and limitations data (i.e. functional group sensitivities and effects of substrate structure on yields, etc.) and separation from the synthesis design program would allow development of a query-driven automatic transform writer. The existence of such a program would do much toward improving the quality and size of the program's database by facilitating and structuring the input of reaction data.

c. Multistep Strategy Implementation in SECS

Currently under development is a SECS facility for allowing the chemist to construct and implement multistep synthetic plans for the target molecule and allow these plans to guide subsequent application of chemical transforms. The plans are updated for each precursor produced in accordance with those long-term goals which have been satisfied by its creation. The planning list structure has been designed such that it is self-modifying, e.g. strategic goals may be activated or removed from the list as called for by other goals in the list, and multiple plans may be implemented simultaneously. The capability for temporally ordering goals is also included so that the chemist may specify the order in which bonds of the target structure are to be disconnected. Design of the module is such that strategic guidance is implemented in the form of both planning goals and strategic constraints. In addition to enabling the use of strategic disconnective and reconnective goals for specifying a general retrosynthetic strategy, the use of constraint information holds promise for shaping synthetic analysis around potential starting materials.

This planning module currently runs in a simulated environment; our plans are to interface it to the synthesis program in the very near future.

d. QED, an Inference Engine for Strategic Chemical Planning

The QED program was developed to provide an experimental workbench for exploring the use of inference and logic in planning chemical synthesis.

QED performs inferences based upon a database consisting of synthetic axioms expressed in the first order Predicate Calculus. The database is first written in mathematical english, and a stand-alone program translates these rules into an associative form useful to QED, as well as allowing for creating new predicates, performing semantic checking, and incremental updating of existing rulebases.

In addition to the separate data base, QED's control structure is based on an agenda list which is created as the rules are examined. Items have priorities assigned when they are created, and QED examines them in the order of their priority. By changing the prioritizing routine, the user can explore the efficiency of various heuristics to guide the search. In this manner we are exploring several approaches, including depth first, best first, and one based on cost-effective examination, where nodes which would not greatly increase the knowledge of the world-picture may be neglected as being too costly to examine.

Currently we are improving QED's database creation routine so that it can examine the rules for improper recursion, and the entire data base for improper recursion, inconsistency and completeness. This will be especially important as more chemists attempt to add rules to the QED data base. We are also allowing QED to create new objects and create a hierarchical inheritance network for these new objects. In this fashion QED could efficiently reason about entities at a higher level with a richer vocabulary, thus extending the power of the inferences.

Finally, we are exploring the axiomatization of the principles of organic synthesis. QED provides a tool for the rapid exploration of various axiomatic schemas with various control structures, thus providing an exciting opportunity to attempt to provide some sort of formal statement of principles governing the design of organic syntheses.

e. Starting Material Strategies - SST

The importance of selecting good starting materials for a synthesis has been known for a long time, but only recently has work started on applying computer techniques to the selection process. We have developed a set of rules useful for the initial selection of potential starting materials (SM's) from an on-line library, and written a program named SST to implement these rules.

The program works in two stages, applying the following rules sequentially:

- 1) Select from the library file the set of molecules whose basic framework is sufficiently close to the target.
- 2) Take the molecules selected and see if their functionality is appropriate for the desired target.

Since questions of functionalization are dealt with at a later point in our algorithm, we did not need all of the fine detail of the molecules in the library. In fact, the high level of detail can hinder a search, obscuring potentially important relationships. We decided to limit the similarity search to carbon graphs, which eliminates the functionality and reduces the size of the search tremendously.

Once the library file and the target molecule have been thus abstracted, we need to apply an algorithm to them to determine which are 'sufficiently close'. We saw four general categories to describe common relationships between target and starting molecules.

- | | |
|-------------------|----------------------|
| I) Target = SM | Identical match |
| II) Target > SM | Superstructure match |
| III) Target < SM | Substructure match |
| IV) None of these | Similarity match |

For a search over our abstracted file, the identical match means that the target and starting materials are identical except for functionalization. The superstructure match is the case where we must make carbon-carbon bonds during a synthesis. The substructure match is the case where the starting material is larger than the target, so carbon-carbon bonds have to be degraded. Finally, the similarity match is where carbon-carbon bonds have to be both made and broken during the synthesis.

The methods for dealing with supergraph and subgraph searching are close but not identical. In particular, for the subgraph search over the abstracted library (Target < SM), the user get the opportunity to further reduce the carbon graph of the target molecule before searching. This might be useful if there are carbon atoms on the target graph which are labile and easily constructed at the end of a synthesis.

Once we have a set of starting material candidates, we then evaluate the functionality to see if it is appropriate.

We have seen good results with this two step algorithm; currently, we are working on extending the first part of it to allow for common subgraph searching over the entire file. In this manner starting materials which need both synthesis and degradation to make the target can be discovered.

2. XENO Program Developments

a. Species Information Implementation

One of the most important biological or hereditary factors affecting the metabolism of a xenobiotic is the metabolic variability between species. Species variations may appear as (1) qualitative differences in the actual pathways of metabolism, and/or (2) quantitative differences in the pathways of metabolism which are common to several species.

In order to represent the quantitative and qualitative species variations in a transform, two types of ALCHEM statements have been developed: One is the qualitative character statement which specifies the species requirement for a transform. This is done in the header of the transform. This character statement, SPECIES <species name(s)>, means the transform is applicable to <SPECIES name(s)> only, not to the other species. The other one is a query statement which reflects the

quantitative variation by adjusting the priority or doing something else for the specified species as shown below:

```
IF SPECIES IS RAT THEN ADD 10
or
IF SPECIES IS MOUSE THEN
  BEGIN
    .
    .
    .
  DONE
ELSE
```

The species protocol for the analysis is established by the user.

b. Biological Activity Evaluation

Many rational methods in structure-activity relations(SAR) studies have been developed. However, these methods have been applied to the SAR of the biological activity and parent compounds rather than metabolites. Frequently, the observed biological activity is not due directly to the xenobiotic compound, instead due to one or more of its metabolites.

We are exploring a unique approach to the problem of evaluating the potential biological activity for metabolites. We have developed a rule-based method for representing the expert knowledge of structure-activity relations. First, we extract the structure requirement (pharmacophore) for the specific biological activity from examples. Then, rules affecting the activity derived from these examples are encoded in the ALCHEM language.

Development of series of rules for each class of compounds to relate structure to biological activity is underway. Current work includes studies on diazo compounds, aromatic amines, polycyclic aromatic hydrocarbons, nitrosamines, and alkylating agents.

c. Collaborative Efforts:

We have been collaborating with metabolism pharmacologists in order to test XENO's prediction for compounds actively being studied in the laboratories. Initial feedback from Mead Johnson on a particular drug indicates that major metabolites were correctly predicted and that some XENO proposed metabolites lead them to further investigation. However, three compounds generated by XENO were not detected in their studies. Other collaborations with investigators at ICI of UK, NIH and the National Australian University are in progress.

D. List of Current Project Publications

- P. Gund, E.J.J. Grabowski, D.R. Hoff, G.H. Smith, J.D. Andose, J.B. Rhodes, and W.T. Wipke, "Computer-Assisted Synthetic Analysis at Merck," J. Chem. Info. and Comput. Sci., 20, 288 (1980).
- S.A. Godleski, P.v.R. Schleyer, E. Osawa, and W. T. Wipke, "The Systematic Prediction of the Most Stable Neutral Hydrocarbon Isomer," Progress in Physical Organic Chemistry, Vol. 13, 1981, pp. 63-118.
- R.E. Carter and W.T. Wipke, "SECS--EH Hjalpmedel Vid Organisk Syntesplanering," Kemisk Tidskrift, No. 7, June 20-25(1981)
- W.T. Wipke and M. Huber, "Symmetry and Organic Synthetic Design", Tetrahedron, submitted.
- W.T. Wipke, G.I. Ouchi, and J.T. Chou, "Computer-Assisted Prediction of Metabolism", in Structure-Activity Correlations as a Predictive Tool in Toxicology, edited by Leon Goldberg, New York: Hemisphere Publishing Corp., 1982 (in press).

E. Funding Status

1. Computer-Assisted Prediction of Xenobiotic Metabolism
Principal Investigator: W. Todd Wipke, Professor, UCSC
Agency: NIH, Environmental Health Sciences
No: ES02845-01
4/1/82-3/31/83 \$ 76,444 TDC
2. Forward Synthesis and Reaction Analysis
Principal Investigator: W. Todd Wipke, Professor, UCSC
Agency: NIH
No: GM 31173-01
6/1/82-5/31/83 \$ 221,381 TDC (Pending approval)

F. Research Environment

At the University of California, Santa Cruz, we have a GT40 and a GT46 graphics terminal connected to the SUMEX-AIM resource by 1200 and 2400 baud leased lines (one leased line supported by SUMEX). We also have a TI745, Heath Data Terminal, CDI-1030, DIABLO 1620, and an ADM-3A terminal used over 300 baud leased lines to SUMEX. UCSC has only a small IBM 370/145, a PDP-11/45, 11/70 and a VAX 11/780, (the 11's are restricted to running small jobs for student time-sharing) all of which are unsuitable for this research. The SECS laboratory is located in a newly renovated room with raised floor in 125 Thimann Laboratories, adjacent to the synthetic organic laboratories at Santa Cruz so the research environment is excellent.

II. INTERACTIONS WITH SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

SECS is available in the GUEST area of SUMEX for casual users, and in the SECS DEMO area for serious collaborators who plan to use a significant amount of time and need to save the synthesis tree generated. Much of the access by others has been through the terminal equipment at Santa Cruz because graphic terminals make it so much more convenient for structure input and output. Demonstrations and sample synthetic analyses were generated for Drs. Terry Brunck and Steve Roman of Shell Development, John Harper of Amoco Chemicals, Prof. Fujiwara, University of Tsukuba, Japan, Dr. Peder Berntsson, Hassle, Sweden. Other visitors included Dr. M. Onozuka, A. Tomonaga and H. Itoh, Kureha Chemical Co., Tokyo, Japan. Demonstrations of SECS in Sweden were performed by Dr. R. E. Carter, University of Lund, Sweden, at many universities and companies. A synthesis of vellerolactone, a substance found to be toxic and teratogenic was generated for Professor Carter. Dr. S. Imahori of Mitsubishi Chemical visited our laboratories for 3 months. Carroll Johnson of Oak Ridge is experimenting with the use of XENO for predicting decomposition pathways in toxic waste materials.

Professor Wipke has also used several SUMEX programs such as CONGEN in his course on Computers and Information Processing in Chemistry. Communication between SECS collaborators is facilitated by using SUMEX message drops, especially when time differences between the U.S. and Europe and Australia makes normal telephone communication difficult. Testing and collaboration on the XENO project with researchers at the NCI depend on having access through SUMEX and TYMNET.

B. Examples of Cross-fertilization with other SUMEX-AIM Projects

This year the SECS and XENO project have made use of the teletype plot program which Ray Carhart of the CONGEN project wrote at Stanford. We modified the program to fit the needs of our projects. This was facilitated by being able to transfer the programs within areas on the same computer system at SUMEX. We continue to have intellectual interactions with the DENDRAL and MOLGEN project in areas where we have common interests and have had people from those projects speak at our group seminars. SUMEX also is used for discussions with others in the area of artificial intelligence on the ARPANET.

C. Critique of Resource Services

We find the SUMEX-AIM network very well human engineered and the staff very friendly and helpful. The SECS project is probably one of the few on the AIM network which must depend exclusively on remote computers, and we have been able to work rather effectively via SUMEX. Basically we have found that SUMEX-AIM provides a productive and scientifically stimulating environment and we are thankful that we are able to access the resource and participate in its activities.

SUMEX-AIM gives us at UCSC, a small university, the advantages of a larger group of colleagues, and interaction with people all over the country. We especially thank SUMEX for support of the leased line for our GT40, and for helping develop our remote print capability.

The only notable area in which the SUMEX facility has fallen short of our needs lies in the occasional periods of persistently high system loading on the KI-10's. In order to obviate this difficulty somewhat, we are currently installing SECS on the 2020 so that we may make use of that system's additional computing capabilities.

D. Collaborations and Medical Use of Programs via
Computers other than SUMEX

SECS 2.9 now resides on the CompuServe computer networks so anyone can access it without having to convert code for their machine. This has proved very useful as a method of getting people to experiment with this new technology. Dr. George Purvis of Battelle is accessing SECS via CompuServe, as are Gene Dougherty of Rohm and Haas and many others. SECS also resides on the Medicindat machine at the University of Gothenborg, Sweden, and is available all over that country by phone. Similarly in Australia, SECS resides at the University of Western Australia and is available throughout the country over CSIRONET. Plans for implementing a similar facility in Japan are currently being examined.

III. RESEARCH PLANS (6/82-6/83)

A. Long Range Project Goals and Plans

The SECS project now consists of two major efforts, computer synthesis and metabolism, the latter being a very young project. Our plans for SECS for the next year include completing the high level reasoning module for proposing strategies and goals, and providing control which continues over several steps. This reasoning module also will be able to trace the derivation of goals and thus explain some of its reasoning. We also plan to focus on bringing the transform library up in sophistication to improve the performance and capabilities of SECS. In particular we plan to allow a transform to have access to the precursors generated as well as the product, this will allow much greater control and more natural transform writing, but it requires extensive changes in the SECS control structure to permit this.

We will continue to explore starting material oriented strategies based on the Aldrich Chemical file we now have implemented. We especially are interested in chirality based strategies which we feel are very strong.

We plan to explore running SECS on a virtual memory 32-bit computer like a VAX-11/780 or a PRIME since many chemistry departments now have these machines available and thus could run SECS.

The XENO metabolism project will be expanding the data base to cover more metabolic transforms, including species differences, sequences of transforms, and stereochemical specificities of enzymatic systems. Development of the second phase which assesses the biological activity of the metabolites will continue as will efforts to simulate excretion and incorporation, the endpoints of metabolism. Finally, application of the

current program to the molecules actively being investigated by metabolism researchers will occur concurrently to test and verify the work done to date on XENO and provide examples for publication.

In the next five years we foresee the SECS and XENO projects reaching a stage of maturity where they will find much application in other research groups. Our research will continue in these areas, but turn to some new programs that approach the problems from different viewpoints and allow us an opportunity to begin fresh taking advantage of what we have learned from the building of SECS and XENO.

B. Justification and Requirements for Continued Use of SUMEX

The SECS and XENO projects require a large interactive time-sharing capability with high level languages and support programs. As a member of the campus computing advisory committee and the UC systemwide computing advisory committee, Professor Wipke has ascertained that the UCSC campus is not likely in the future to be able to provide this kind of resource. Further there does not appear to be in the offing anywhere in the UC system a computer which would be able to offer the capabilities we need. Thus from a practical standpoint, the SECS and XENO projects still need access to SUMEX for survival.

Scientifically, interaction with the SUMEX community is still extremely important to our research and will continue to be so because of the direction and orientation of our projects. Collaborations on the metabolism project and the synthesis project need the networking capability of SUMEX-AIM, for we are and will continue to be interacting with synthetic chemists at distant sites and metabolism experts at the National Cancer Institute. Our requirements are for good support of FORTRAN.

C. Needs Beyond SUMEX-AIM

We do plan to acquire a virtual memory minicomputer like a VAX or PRIME in the future to offload some of our processing from SUMEX. Such a machine would enable us to do some production and development work locally and would explore the feasibility of those types of machines as hosts for SECS and XENO. Acquisition of a local machine would lead to decreased loading on SUMEX, although we anticipate a continued need for SUMEX access since (1) we plan to continue to develop and maintain the PDP-10 version of SECS and (2) we require the networking capabilities of SUMEX for purposes of collaboration and testing. In the event that no local facilities are forthcoming, we foresee our SUMEX load contribution increasing as our group grows and as we start new projects and maintain existing large programs.

D. Recommendations for Community and Resource Development

The AIM Workshops have been excellent in the past and should be continued. In view of current congestion of the SUMEX resource, we feel the community would benefit if large remote users such as ourselves had a virtual minicomputer in order to distribute the computing load more effectively. This would free SUMEX, with its networking facilities, to

allow netwide testing/utilization of SECS and XENO by members of the community and by groups working in collaboration with our own.

II.A.2.7 SOLVER Project

SOLVER: Problem Solving Expertise

Dr. P. E. Johnson
Center for Research in Human Learning
University of Minnesota

Dr. W. B. Thompson
Department of Computer Science
University of Minnesota

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

This project focuses upon the development of strategies for discovering and documenting the knowledge and skill of expert problem solvers. In the last fifteen years, great progress has been made in synthesizing the expertise required for solving extremely complex problems. Computer programs exist with competency comparable to human experts in diverse areas ranging from the analysis of mass spectrograms and nuclear magnetic resonance [DENDRAL] to the diagnosis of certain infectious diseases [MYCIN].

Design of an expert system for a particular task domain usually involves the interaction of two distinct groups of individuals, "knowledge engineers," who are primarily concerned with the specification and implementation of formal problem solving techniques, and "experts" (in the relevant problem area) who provide factual and heuristic information of use for the problem solving task under consideration. Typically, the knowledge engineer, after consulting with one or more experts, decides on a particular knowledge representational structure and inference strategy. Next, "units" of factual information are specified. That is, properties of the problem domain are decomposed into a set of manageable elements suitable for processing by the inference operations. Once this organization has been established, major efforts are required to refine representations and acquire factual knowledge organized in an appropriate form. Major research problems exist in developing more effective representations, improving the inference process, and in finding better means of acquiring information from either experts or the problem area itself.

Programs currently exist for empirical investigation of some of these questions for a particular problem domain [AGE, UNITS, RLL]. These tools allow the investigation of alternate organizations, inference strategies, and rules bases in an efficient manner. What is still lacking, however, is a theoretical framework capable of reducing dependence on the expert's intuition or on near exhaustive testing of possible organizations. Despite their successes, there seems to be a consensus that expert systems could be

better than they are. Most expert systems embody only the limited amount of expertise that individuals are able to report in a particular, constrained language (e.g. production rules). If current systems are approximately as good as human experts, given that they represent only a portion of what individual human experts know, then improvement in the "knowledge capturing" process should lead to systems with considerably better performance.

B. Medical Relevance and Collaboration

Collaboration with Dr. James Moller in the Department of Pediatrics, Dr. Donald Connelly in the Department of Laboratory Medicine, and students and staff in the University of Minnesota Medical School.

C. Highlights of Research Progress

Accomplishments of this past year. Prior research at Minnesota on expertise in diagnosis of congenital heart disease has resulted in a theory of diagnosis and an embodiment of that theory in the form of a computer simulation model which diagnoses cases of congenital heart disease. This past year the work has been extended to a system called Galen.

Galen is a computer program that simulates the diagnostic process in pediatric cardiology. It is descended from two earlier programs written here at Minnesota: Diagnoser and Deducer [Swanson, 1977]. Deducer is a program that builds hemodynamic models of the circulatory system that describe specific diseases. The models are built by using knowledge about how idealized parts of the circulatory system are causally related. Diagnoser is a recognition-driven program that performs diagnoses by successively hypothesizing one or more of these models, matching them against patient data. The models that match the best are used as the final diagnosis. A series of experiments carried out at Minnesota have shown that Diagnoser/Deducer performs as well (and sometimes better) than expert human cardiologists [Johnson et al., 1981].

Despite the success of Diagnoser and Deducer, the programs have become increasingly difficult to use as a research tool in recent years. They lack a clean, comprehensible structure that is necessary for the kind of experiments we wish to perform. To remedy this problem, a new version, called Galen, is nearing completion here. Galen is intended as a version of Diagnoser/Deducer that is easy to use and modify.

Like its predecessors, Galen is strongly recognition-driven. The program consists of three knowledge sources, referred to as the guesser, the reviewer, and the modeler. The guesser is a set of production rules that propose initial hypotheses on the basis of patient data: a "first guess" about what is wrong with the patient. A hypothesis can be regarded as a set of models (or other hypotheses) and a set of rules that choose the most appropriate members of the set for further consideration. The reviewer uses these hypothesis-specific rules to refine each hypothesis to the most specific form that still satisfactorily accounts for the data. The modeler is a group of production rules that builds new models and

attaches them to existing (precompiled) hypotheses. It runs in isolation from the other two knowledge sources, and is used to initialize the system.

In a typical diagnostic session, Galen is given abstracted data one piece at a time from the patient's chart. Each new piece of data is examined by the guesser. Next, each hypothesis that is currently in contention is examined by the reviewer. When this is complete, a new piece of data is requested and the process repeats until all data has been examined. Once this occurs, the models that remain in contention are scored according to how well they explain the data. The two or three top scoring models are then used as a final diagnosis.

We are investigating several research questions within the architecture of the Galen system. First, we are experimenting with ways of qualitatively scoring cardiovascular models, to replace the somewhat ad-hoc numeric scoring method now in use. Second, we are investigating ways in which Galen can be made to ask for additional information about the patient at the appropriate time. Third, we are studying ways in which the modeler (our causal reasoning component) can be integrated with the guesser and the reviewer (our prototypic reasoning components). In particular, we are interested in exploring ways in which causal reasoning can replace (or aid) prototypic reasoning when it proves inadequate to reach a diagnosis.

Research in progress. The methodology of our research derives from the discipline of cognitive science, and from our study of expert problem solvers in a number of fields. This methodology consists of: (1) extensive use of verbal thinking aloud protocols as well as other experimental data as a source of information from which to make inferences about underlying cognitive structures and processes; (2) development of computer models as a means of testing the adequacy of inferences derived from the protocol studies; (3) testing and refinement of the cognitive models based upon the study of human and model performance in experimental settings.

We are constructing recognition based computer models of diagnosis in several fields (medicine, law, science) with expertise comparable to highly expert practitioners. Since humans are notoriously poor at describing their own perceptual knowledge, the first phase of this work requires creating a series of tasks through which experts can reveal their criteria for initiating hypotheses and the associated expectations required for acceptance or rejection. The first type of task requires experts to solve real and simulated problems. Based on studies of performance in these tasks, a model is created to explain the ways in which recognition prototypes are selected for further evaluation. A catalog of these prototypes are produced. Expert performance cannot, however, be realized until these prototypes are associated with highly accurate recognition templates to be matched against input data. While our experts are not particularly adept at generating these templates, we have discovered they can frequently describe causal models of the underlying process. We therefore attack the problem of generating templates by first constructing a causal model of relevant parts of the underlying process. The model is intended as a tool to "capture" knowledge useful for a non-causal reasoning system. In a sense, we use a causal model to "compile" a recognition model